The NIST Display Interface Testbed: the Development of a Centrally Controlled Testing System for the Automated Generation of Stimuli and Control Feedback

John W. Roberts and James A. Ward

National Institute of Standards and Technology[1], MS 8951, Gaithersburg, MD 20899

ABSTRACT

Impressive advances are being made in the field of display metrology, as illustrated in standards such as the Video Electronics Standards Association's Flat Panel Display Measurement standard. However, issues remain regarding the technology-dependent response of displays to large-scale signal errors such as noise-induced errors on the data and control lines. The NIST Display Interface Testbed can be used in conjunction with a graphics controller, test generator, or other standard signal source, to inject a user-specified controlled sequence of errors into the interface of a display under test, permitting observation and measurement of the response of the display. In addition to its function as an interface tester, the Testbed effectively serves as a non-intrusive probe into the operation of the display. The programmed control structure of the Testbed permits the convenience of automated test sequences. Development of a feedback mechanism which allows the test stimulus to be modified in real time, as a function of measurement results and user commands is planned for 1999.

keywords: display interface, display measurement, noise generation, test automation

1. INTRODUCTION

Like all electronic devices, display systems must somehow function in an imperfect world. Logic and timing errors in the signal path of a display make up one of the classes of imperfections that must be considered in the design or selection of a display. A controversy regarding the "acceptable" level of errors in a display signal interface led to an interest in finding a way to determine the response of a display to various bit error rates. This was the inspiration for the development of the NIST Display Interface Testbed, or "DIT".

The Display Interface Testbed was developed to test the responses of displays to errors (on the time scale of one or more pixel clocks) in the data and control signal paths. The DIT takes a normal signal from a computer-driven graphics controller or commercial test signal generator, then injects a user-specified artificial noise signal into the signal path. The impact of this noise signal on display performance can be observed visually, or measured quantitatively using methods such as those described in the Video Electronics Standards Association (VESA) display measurement standard. In addition to the usual static measurements, attention should be paid to transient phenomena that will be associated with intermittent signal errors on fast-response displays.

The current working prototype Display Interface Testbed is set up for use with a digital display interface using TTL voltage level signals - a type of interface still used by a considerable number of flat panel displays. A modification currently being designed will allow the DIT to use the 15-pin analog "VGA" interface found on most current CRT-based computer monitors as well as many commercial products incorporating newer technologies, such as LCD desktop monitors, plasma display panels, and LCD or DMD (digital micromirror) projectors. An implementation is planned for the future to allow many HDTV (high definition television) displays to be tested.

In addition to its original objective of permitting the evaluation of displays driven by degraded signals, the DIT offers several other advantages for display measurement. The modular construction of the hardware and control software allows straightforward reconfiguration for use with new classes of displays and new display interfaces. New test functions can also be added readily. These new functions include the ability to partially automate the test process, which reduces testing time and permits more complex test sequences to be employed.

---

[1] Certain commercial products are identified here in order to document our experiments. Identification of such products does not imply endorsement by NIST.

The long-term objective of this project is to establish the utility of the DIT as an aid to display measurement and evaluation. The technology developed by the project will be made available to industry for adaptation and incorporation in future test devices and testbeds.

This paper discusses the need for the types of measurements made possible by the DIT (including the use of induced errors at the scale of the pixel clock), the principles of measurement using the DIT, a more detailed description of the DIT architecture and operation, a discussion of how measurements can be made with the aid of the DIT, and several ways in which the DIT can be configured and programmed to automate the testing process.

2. THE NEED FOR PIXEL-SCALE SIGNAL-BASED DISPLAY PERFORMANCE MEASUREMENT

The measurements made possible or facilitated by the Display Interface Testbed can supplement existing measurement methods, permitting more accurate characterization of displays and in many cases reducing the time and effort needed to evaluate displays. Different benefits should apply to different types of users:

- Display manufacturers can incorporate DIT-style tests during the design process to make sure a new display meets design criteria in terms of response to signal errors, and after design is complete while writing the signal specifications for the display.
- Systems integrators can use the DIT to evaluate candidate displays for their system designs, to select displays that can tolerate the environmental conditions they expect, and to determine how tightly they have to control the operating environment of a particular display in their application.
- Any users who need to operate a display in a specific operating environment can use the DIT to compare the performance of different displays in that environment. A user who is aware that the operating environment requirements may change over time should use an implementation of the Display Interface Testbed that includes the easily modifiable architecture of the original.
- Certain users may need a display for use in applications that were not originally anticipated by the designers of the displays, and for which the performance of the displays is not known. Knowledge of the internal operation of the displays (which may be needed to predict their response in these applications) may be held as proprietary information by the display manufacturer. If the anticipated market is small, the user may be unable to persuade a display manufacturer to design a display for a specific application. In such a situation the DIT may be able to emulate the conditions of the desired application, so that different existing displays can be compared for that specific application.

3. THE DISPLAY INTERFACE TESTBED AND THE PRINCIPLES OF MEASUREMENT

An important step in planning the design of the Display Interface Testbed is the development of theoretical models for display, signal noise, and the response of the display to a given signal stimulus. The operating parameters of the DIT can then be set in accordance with the model, and experiments performed to determine how well the models match reality. The modular structure and reconfigurability of the DIT permit it to be modified to match changes in the models. Once the all the models are developed, a set of operational rules must be established to control the way in which the measurement process is conducted.

The display assembly can be regarded as the display itself, plus the display interface, which typically transfers both image data and control signals (sync signals, pixel clock, etc.). Errors in the transferred image data will result in errors in the image produced by the display device. Errors in the control signals will produce results that are highly dependent on the details of the internal display electronics, and therefore difficult to predict without actual tests. For a variety of reasons, including frequency of transitions, it is also expected that control signal errors will be considerably less frequent than image data errors in a real-world display interface. The simplest display model, therefore, is one in which all signal errors on the image data inputs result in visible defects in the image, and the control inputs are not subject to errors. While not fully realistic, it is felt that this display model is adequate for the initial design of the DIT, and that testing of response to errors in control signals can be performed on a case-by-case basis as desired.

The simple display model can be improved by adding known characteristics of a display to be tested, that are likely to have an effect on the display performance as the DIT test stimulus is applied. For example:

- Active matrix LCD (AMLCD) and field emission (FED) displays usually exhibit fast pixel response time, which improves their usability for video (moving images), but which also tends to make single-pixel errors particularly visible. Passive matrix LCD (PMLCD) displays, in comparison, often have much longer response time, which has the disadvantage of smearing moving images, but which also greatly reduces the visibility of single-pixel errors - it is believed that even a relatively high rate of random errors in image data would appear mainly as a reduction in contrast. Non-random errors could produce localized artifacts on a PMLCD.

- Plasma and DMD technologies (among others) combine relatively fast pixel response and lack of inherent grayscales. At any given instant, any particular pixel is either fully "on" or fully "off". Intermediate intensity levels must therefore be implemented by temporal modulation, spatial modulation, or a combination of the two. DMD temporal modulation is typically performed within the time of a single frame period (this is made possible by the extremely fast switching rate of the micromirrors), so image errors are unlikely to be propagated from one frame to the next. Plasma displays may extend temporal modulation over a number of frames, so input pixel errors that enter the panel during a particular frame period may affect the appearance of neighboring frame periods as well. Similarly, spatial modulation may cause the effects of a single pixel error (or a short burst of pixel errors) to be spread out over a larger area of the display.
- Many cathode ray tube (CRT) displays do not exhibit a fixed relationship between the logical pixels defined by the image data and the phosphor elements on the screen. Most CRT displays intended for use as computer monitors can be operated in multiple resolution modes, with different mappings for each between logical pixel and phosphor element. The electron beam that scans the phosphors is usually wider than the "dot pitch" (distance between corresponding phosphors), but with intensity greatest near the center of the beam. As a result, the effects of a "single pixel" error tend to be spread over an area of the screen larger than a single logical pixel.

The noise model can start with a simple memoryless temporal distribution of induced errors, and the assumption that an "error" represents a value different from that of the original input value from the signal source. As an example for digital display interfaces, the data for an error pixel could be represented by the one's complement of the original value. Other models for the value of an error pixel could call for it to be set to the minimum value or the maximum value of the permitted intensity range, which could be useful for some experiments. Timing other than exponentially distributed interarrival times can also be implemented. For example, an effort can be made to model "bursty" noise patterns. The overall pixel error rate can be set at a level that might be encountered in a working system, or it may be set to an artificially high level in order to evaluate a display over a wider range of stimulus inputs.

The operating procedures for the Display Interface Testbed should allow for useful results from a nondestructive testing procedure:

- The test image produced by the signal source should be selected so that induced errors may be readily detected. In most cases, relatively large areas of uniform color will make individual pixel errors stand out more clearly than with a test image using a highly complex pixel-scale pattern. A mostly light test image can be useful if error pixels are defined as "dark". It should be noted that these are merely guidelines, in the absence of specific requirements for a particular measurement. Some standard measurements call for a specific pattern, which must therefore be used even if it means that the detection of error pixels will be made more difficult by the use of that pattern. It should also be noted that some measurements pertaining to the operation of video systems may require a test pattern of moving images.
- The noise signal is ordinarily chosen to be "representative" of real-world noise characteristics, whether an exact reproduction or deliberately exaggerated by some factor. An exception to this guideline would be appropriate when the DIT is deliberately being used as a probe into the internal operation of a display, and the stimulus is not constrained to be anything that a display would encounter in actual use.
- Induced errors in control signals (as opposed to image data signals) require special attention - for example, care should be taken that the display is not damaged by an unusually high level of such errors, or by a sequence that causes the display electronics to "lock up".
- In most applications, it is desirable to perform the DIT tests without causing permanent damage to the display being tested. In addition to the risk of damage from errors in control signals, some displays could potentially be harmed by certain patterns of errors in the image data. For example, LCD displays should avoid periodicities in the noise signal that closely match internal timings such as the frame rate, to prevent the risk of a DC bias being applied to some of the liquid crystal material. If there is a periodicity in the noise signal, it should be chosen to be "relatively prime" to likely internal timings of the display being tested. Note that damage to a display, when subjected to an exact reproduction of expected operating conditions, may indicate either that the display being evaluated should not be used for that application, or that the operating conditions were not reproduced with sufficient accuracy.

4. HARDWARE ARCHITECTURE

The current working version of the Display Interface Testbed makes use of an arbitrary waveform generator (AWG) with digital as well as analog outputs, to generate the noise signal for a digital-input display. The model of AWG used has 32 digital outputs, so the noise on up to 32 signal lines can be controlled separately, or multiple lines can be controlled together by rewiring the interface hardware. The AWG can be set to use the pixel clock from the signal source in place of its internal clock, so the noise signal is generated on a pixel by pixel basis. The signal source being used at the moment is a personal computer with a digital-output graphics controller card. The DIT interface hardware injects the synthetic noise signal into the normal signal path of the display interface. The noise model used in the design of the current version of the DIT specifies a noise value as a logical inversion (or one's complement in the case of a multiple bit data value) of the correct signal, so in this case the combination of display signals and noise signals can be accomplished by a bank of exclusive-OR (XOR) logic gates, with the outputs connected to the display, one input of each gate connected to a line from the display signal source, and the other input of each gate connected to a line from the noise generator. If the signal delay caused by the noise generation and injection process were to cause an unacceptable skew among the signals used by the display, a time delay could be introduced in the circuitry to minimize the skew. This has not been necessary in the versions implemented thus far.

The computer that controls the DIT operates the AWG through an IEEE-488 (GPIB) instrument interface. Real-time control of the signal source (to allow changing the test image during the course of a test sequence) has not yet been implemented, but this can be done using GPIB or serial interface (in the case of a test generator), or by network connection (in a configuration like the present one, where a separate computer generates the display signal). Figure 1 is a generic block diagram of the current version of the Display Interface Testbed.

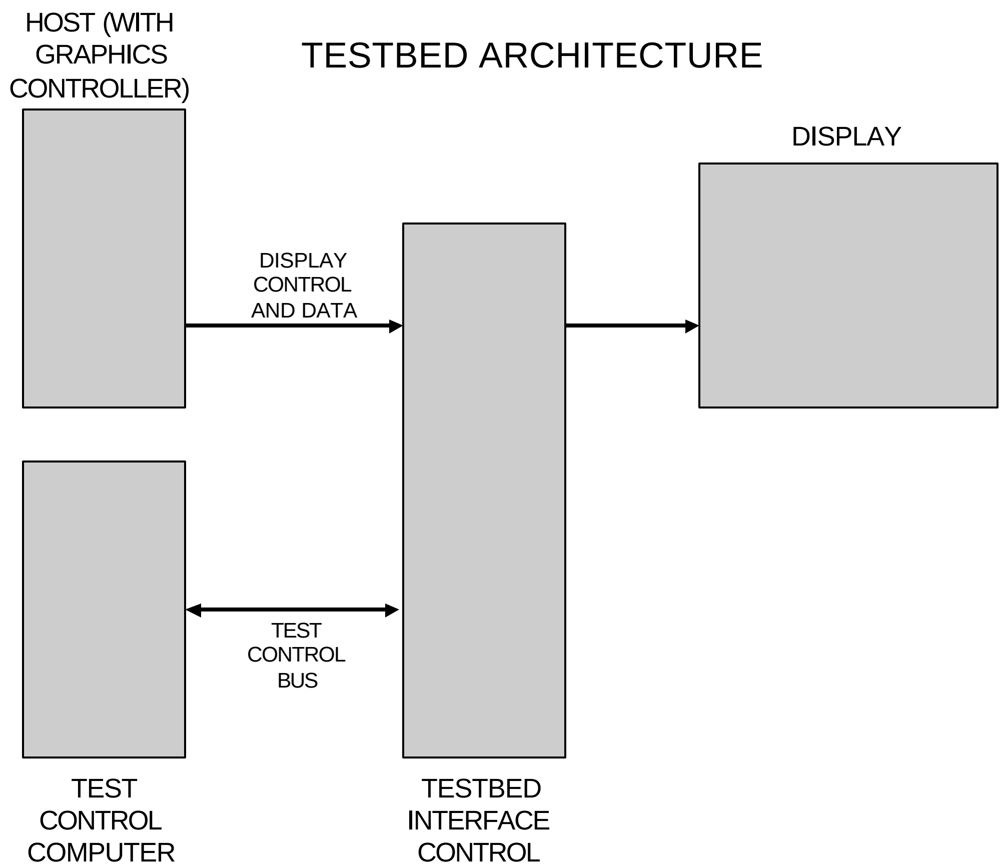

Figure 1. Generic block diagram of the Display Interface Testbed

Measurement of the display under the stimulus of the DIT can be performed using standard display measurement equipment, independent of the operation of the DIT. Modifications are being developed that will also permit direct involvement of the DIT in the measurement process, including:

- Data logging,
- Automated control of complex test sequences, and
- Feedback from display measurements, to guide the course of a test sequence.

5. SOFTWARE ARCHITECTURE

The Display Interface Testbed controller is programmed in LabVIEW(R), a commercial general-purpose programming language which provides a graphical programming and user interface, plus a library of routines to facilitate communications and control of laboratory instruments via GPIB interface, parallel and serial ports, network connection, etc. The fundamental organizational unit of a LabVIEW program is the Virtual Instrument (VI), which largely corresponds to a routine or subroutine in other programming languages. A VI can be a "sub-VI" of another VI, and in turn can contain sub-VIs of its own. LabVIEW is a data-driven language; processes (VIs) stand ready to execute at all times, waiting for all necessary data inputs to arrive. A VI may contain loops, conditional tests, absolute numeric values, etc. Laboratory instrument manufacturers often provide VIs written to allow direct control of their instruments by LabVIEW. LabVIEW programs are often written for real-time operation, and for real-time interaction with a human user.

The current version of the DIT uses a repeating pseudorandom sequence stored in the arbitrary waveform generator as the synthetic noise source. This sequence is stored as a "noise file"; the number of elements in this file (and thus the periodocity of the noise sequence) must be relatively prime to any of the major periodicities of the display (for example, the number of clock cycles per frame), if the noise source is to serve reasonably well in the place of a truly random sequence. The data values stored in the noise file are typically read out one entry per pixel clock cycle, and used to affect the pixel(s) being addressed during that clock cycle. The original display signal and the noise signal are combined using an array of XOR logic gates, so a zero bit in the noise file means that the data on one signal line will be passed to the display unchanged during that clock cycle, while a one bit in the noise file means that the corresponding signal line value will be inverted during that cycle. If all of the bits corresponding to data lines are programmed to go high or low together for each clock cycle (in other words, if all of the values in the noise file are either 0000... or 1111...), then the "noise" pixels on the display will appear to be the inverse of what they should have been - a pixel that was supposed to be dark will appear light, a pixel that was supposed to be light will appear dark, colors will be reversed, and so on. Alternatively, individual signal lines or groups of lines can be controlled independently. For example, the red subpixels can be subjected to noise while the green and blue subpixels are unaffected, or only some of the bits (for example, the three least significant bits) per color can be affected. The next version of the signal combination logic for the DIT will allow for more complex noise models – pixels can be inverted, driven to maximum brightness, or driven to minimum brightness. The modular structure of the software allows this hardware change to be accommodated by changing just one module (VI) of the control software.

The software of the current version of the DIT consists of a front panel for user interaction, a generator to create the noise file, a utility to transfer the noise file to the AWG, and a utility to initialize the AWG and start output of the noise signal. The front panel is the main VI for the DIT program, and the other functions are sub-VIs accessed through the front panel. VIs can be replaced or added as the the DIT hardware evolves and as new tests are needed.

The front panel in the DIT LabVIEW program displays a number of control buttons - once the program is running, clicking on these buttons activates the corresponding sub-VI.

The noise generator sub-VI opens a window and prompts the user for the desired error rate, number of entries in the noise file, and a filename (the AWG can hold multiple noise files in its memory, thus permitting rapid switching between noise files to change the characteristics of the noise signal being produced). The user-specified error rate is used to compute a threshold value. During the generation of the elements for the noise file, a standard pseudorandom generator produces a sequence of random numbers, which are compared to the threshold value. Each corresponding entry in the noise file is set to either the "noise" value (1) or the "no noise present" value (0), based on the results of this comparison. Given a sufficiently large noise file, the effective noise level can therefore be set very close to the desired value, with mathematical characteristics similar to those of at least one type of natural noise.

The file transfer sub-VI converts the data in the noise file into the format used in GPIB communications, initiates a connection to the AWG, transfers the data, then sends a number of device-specific commands that instruct the AWG to store the data as a signal file, and that give the user-specified filename. (If the user did not specify a filename, a default name is assigned.) The sub-VI to start the test process does so by sending a sequence of commands to the AWG, specifying the signal (noise) file to be used and the AWG configuration settings, and finally giving the command to begin output of the signals in the noise file. Additional commands accessed by the front panel halt the operation of the AWG when the measurement process is completed, allowing the user to enter new noise files, and exit the LabVIEW program when the sequence of tests is completed.

6. OPERATION OF THE DIT (CURRENT CONFIGURATION)

The first step in operating the Display Interface Testbed is to connect the signal source (computer with graphics controller, test generator, etc.), the DIT interface logic, and the display to be tested. At present there are many different digital display interfaces in use, so adaptation for a new display requires considerable custom configuration. Work is underway to simplify the connection process by making the circuitry easier to reconfigure. At the same time the development of digital display interface standards promises eventually to reduce the extent to which reconfiguration will be necessary. The analog "VGA" adapter now being developed will work with a wide range of displays with no further configuration needed.

The second step is to study any information that may be available regarding the display to be tested (e.g. horizontal and vertical addressability, timing, and the number of clock cycles per line and per frame). This information will allow the user to choose the number of entries in the noise file (to make sure it is relatively prime with respect to panel timings), and to decide what sorts of tests are desired. The DIT can be operated without the benefit of this information, but greater caution is needed to avoid the risk of damage to the display.

Once the components have been connected and the parameters of the test established, the user starts up the control computer, turns on the signal source and AWG, turns on the display (making sure that any power sequencing requirements are met), and starts the DIT control program. In a typical case, the user would order the creation of a new noise file meeting a new set of specifications, transfer the file to the AWG, then run the actual test.

Even in cases where precision measurements are to be made by instrument, visual observation can often be used to judge the correct operation of the DIT. For example, if the noise is specified to be random, but motionless or slowly moving patterns of noise pixels appear on the display, then the periodicity of the noise signal is inconveniently close to some periodicity of the display. In this case, a different size noise file should be specified.

It is usually desirable to avoid damage to the display, and the user should bear in mind that certain types of noise signals could potentially cause degradation of some types of display over a long period of time (e.g. by the introduction of a DC bias to the material of a liquid crystal display). The current DIT has been operated for minutes at a time (with a cumulative time of many hours) on a liquid crystal display without any apparent damage to the display, but until more is known it is advisable to avoid running the DIT unattended for hours at a time, especially at a high noise level. Fortunately, automated testing should substantially reduce the cumulative run time needed to evaluate a particular display. A working implementation for industrial use will have safety limits designed into the signal settings to prevent accidental damage to displays.

7. MEASUREMENT TECHNIQUES USING THE DISPLAY INTERFACE TESTBED

Evaluation of displays using the DIT has thus far been restricted to visual observation. Another group at NIST within the Electronics and Electrical Engineering Laboratory (EEEL) has considerable experience performing the photometric and colorimetric measurement of displays The group within the Information Technology Laboratory (ITL) that has developed the DIT feels that it would not be productive to attempt to duplicate such a capability, or to perform the pioneering work in development of such measurement techniques that the EEEL group has undertaken. Plans are being discussed for a joint effort. As an example, the EEEL group can provide the expertise needed to implement precision measures with use of the DIT as a stimulus, and the ITL group can make the DIT available to the EEEL group for use in performing such measurements.

As a general rule, precision measurements of displays being driven by the DIT would be useful for obtaining information on display performance in the presence of degraded input signals. Initial measurement techniques are likely to be similar to "regular" display measurement techniques, such as those described in the VESA display measurement standard. Measurements such as contrast ratio, taken over a wide range of noise levels, may be able to provide insight into the internal operation of the display, and possibly allow some form of extrapolation to performance under noise levels that are too low to register degradation using current measurement techniques. Refinements of the measurement process are likely to emphasize detections of transients - pixel defects and display artifacts that appear and disappear on a frame-by-frame basis. Some of the performance parameters that the DIT can provide as stimulus, such as single-frame, single-pixel defects, may require development of new measurement techniques.

Future versions of the DIT may include the capability to reliably force a pixel defect at a specific location on the display. This will require highly reliable triggering on sync signals and pixel clock counting to avoid "jitter" that would cause the defect to move around slightly. The prior knowledge of of an upcoming bad pixel in a known location could assist substantially in the development of new techniques for measurement of transient display phenomena.

8. EXPERIMENT CONTROL: GENERAL ISSUES

Many experiments are conducted using the classical methods of setup entirely by hand and data recording by hand. This is usually the simplest method, and often the quickest way to obtain data from one test setup. However, it severely limits the ability to collect large amounts of data, and hinders changes in the parameters of the experiment. With care, partial automation can be added to a test system, without too great an expense to the user in time and money, providing substantial improvements in data collection and reconfigurability. Initial refinements include the ability to perform automated test sequences and automated data collection.

9. USING THE DIT TO CONTROL AN EXPERIMENT

The current version of the Display Interface Testbed provides a function menu on the front panel which guides the user through a sequence of steps needed to conduct a measurement session: generation of a "signal noise file", transfer of the file to the AWG, initiation of AWG output to begin the experiment, and termination of the experiment. As components are added and the capability of the DIT increases, the control sequence will become increasingly complex. The user will not necessarily want to operate each control every time a test is run. The next step in automation is therefore to program default test sequences, which the user can select to cause the DIT to automatically run through a number of steps in the test sequence. Full step-at-a-time control should still be permitted, but automated sequences should be suitable for a large percentage of test situations, and will substantially reduce the time required to run an experiment. As an example of an automated sequence, the user could specify an initial set of test conditions, then command the DIT to run a test n times, with a different noise level for each run, and a single prompt from the user to move on to the next step.

The incorporation of measurement instruments under LabVIEW control is expected for future versions of the DIT. As a simple example, a high speed digital camera can be attached, with image analysis software to determine the distribution of artifacts on the display. Some commercial photometric/colorimetric instruments also provide for computer control and readout of data. Under program control, the DIT will be able to store these measurement results in memory, and in some cases perform partial analysis of the data in real time.

10. AUTOMATION AND FEEDBACK

A further refinement for automated test sequences is feedback from measured results to adjust the course of the test sequence. Use of this capability offers the possibility for massive speedup of certain types of tests. For example, suppose a certain type of display experiences degraded contrast ratio in the presence of noise on the display interface. Further suppose that the user wishes to know the noise level that will produce a specified contrast ratio. Rather than running an extended sequence of tests with predefined noise levels and plotting a graph to determine the threshold value, the user can instruct the test instrument to adjust the noise level until the specified contrast ratio is observed. If desired, the intermediate measurements can be recorded and evaluated as a check on unexpected phenomena, but the point is that the target threshold number can be obtained quickly and automatically. The speedup in the measurement process is even greater if two or more input parameters are to be adjusted. A substantial speedup not only facilitates laboratory measurement, but also offers possible application in a manufacturing environment. Figure 2 shows an example of an advanced configuration with a feedback path from one type of measuring device.

The use of control automation and feedback is certainly not new, but the design approach of the Display Interface Testbed makes it particularly convenient to add these capabilities for use in measurement sequences involving degraded control signals.

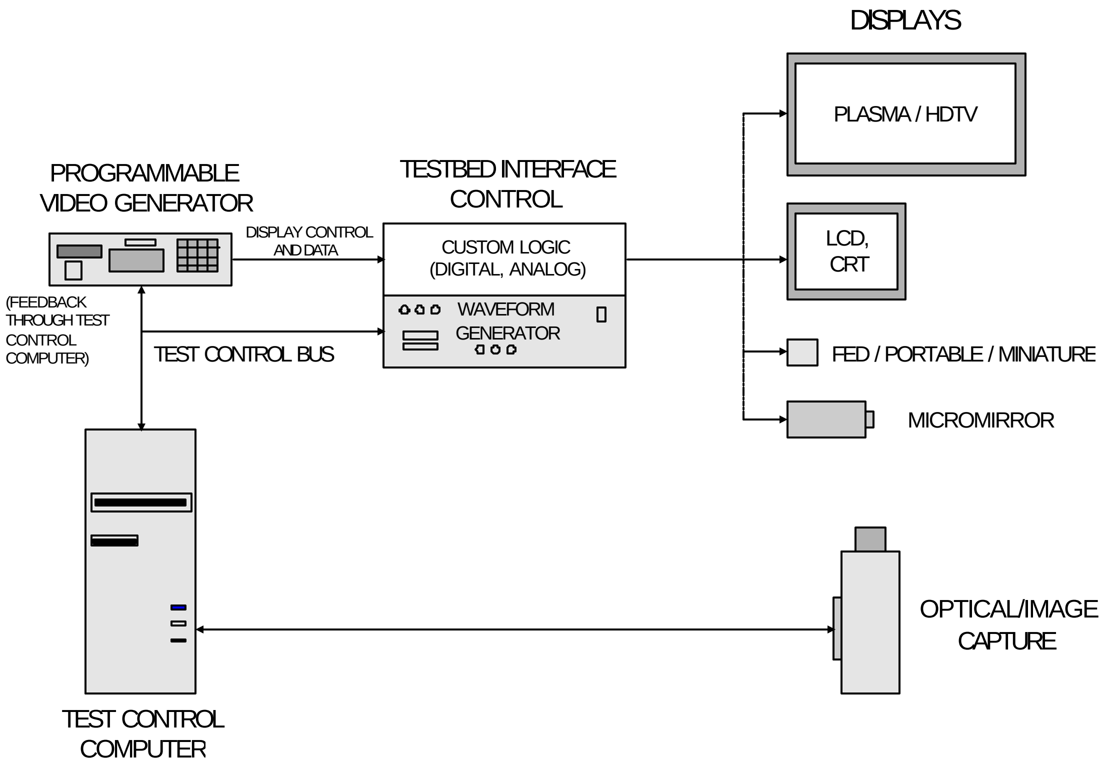

Figure 2. Example of advanced DIT configuration using measurement feedback

11. PLANS FOR THE DISPLAY INTERFACE TESTBED

A number of steps are either being implemented now or planned for future improvement of the Display Interface Testbed:

- Develop a hardware signal interface that simplifies switchover from one test display to another.
- Complete development of one or more interfaces specifically for use with analog-input displays.
- Add the ability to use with low voltage differential interface protocols.
- In addition to the signal noise model currently being used, offer the user a choice of noise stimuli with different mathematical properties.
- Add the ability to generate noise pixels at specified locations on the screen (temporally random, but spatially nonrandom).
- Add a control path to permit the DIT software to change the test image in real time.
- Increase the use of automated test sequences, as DIT funtions grow.
- Look for instruments that allow automated readout of data, and provide for data logging (plus some level of real-time data analysis).
- Plan for the eventual incorporation of feedback control to adjust the course of an experiment.
- Construct specialized hardware to generate the noise/stimulus signal, replacing the function of the AWG and permitting use with the highest resolution displays.
- Work with EEEL on measurement techniques, initially on adding simple measurement capability, and eventually on resolving issues regarding measurement of transient phenomena and automated capture of data.

12.SUMMARY

The Display Interface Testbed was developed as a tool to facilitate evaluation and measurement of displays operating with errors in their input signals. Conceptual models of the components of the measurement system have been developed, and continue to be refined. A basic working version of the DIT has been constructed and used with visual evaluation of the results of the error signals on the display output. Existing measurement techniques can be used with the DIT, and there are plans being developed to add integrated measurement capability, in cooperation with the NIST EEEL display measurement project. Transient phenomena (both error pixels and display artifacts specific to particular display technologies) are expected to be of particular interest in measuring displays with noisy input signals.

In addition to its particular usefulness for noise-related measurements, the modular construction, reconfigurability, and program control of the DIT facilitate the measurement process in a number of ways. For example, they simplify the addition of test sequence automation, measurement data logging, and feedback control of experiments.

The Display Interface Testbed has been shown to work, and it makes possible or facilitates new categories of display performance measurements that can be used to predict the performance of a display in a "real world" application. Continued improvements will enhance the capability and usability of the DIT for this purpose.

13. ACKNOWLEDGEMENTS

This work was supported by the Advanced Technology Program (ATP) at NIST. (contact http://www.atp.nist.gov/)

14. REFERENCES

1. A. Donohoe, "Display Image Tolerance Testbed", unpublished NIST internal document, 1998.
2. VESA Flat Panel Display Measurements Standard, Video Electronics Standards Association, 1998.
3. "LabVIEW(R) User Manual", National Instruments Corporation, Austin, TX, 1998.